

Cabling's contribution to energy efficiency



Rob Cardigan
Nexans Cabling Solutions



“By 2020
IT will be more
polluting than
aviation”

(the Economist)

“**Unrestrained** IT power consumption could eclipse hardware **costs** and put great pressure on **affordability**, data center infrastructure and the **environment.**”

Luiz André Barroso, (Google)
“The Price of Performance”



Typical 3000m² Data Centre

- † Consumes ~ 45MWhr/yr
- † Equates to 24,000 tons CO₂
- † Equivalent to ~1600 Cars



Data Centre IT Load

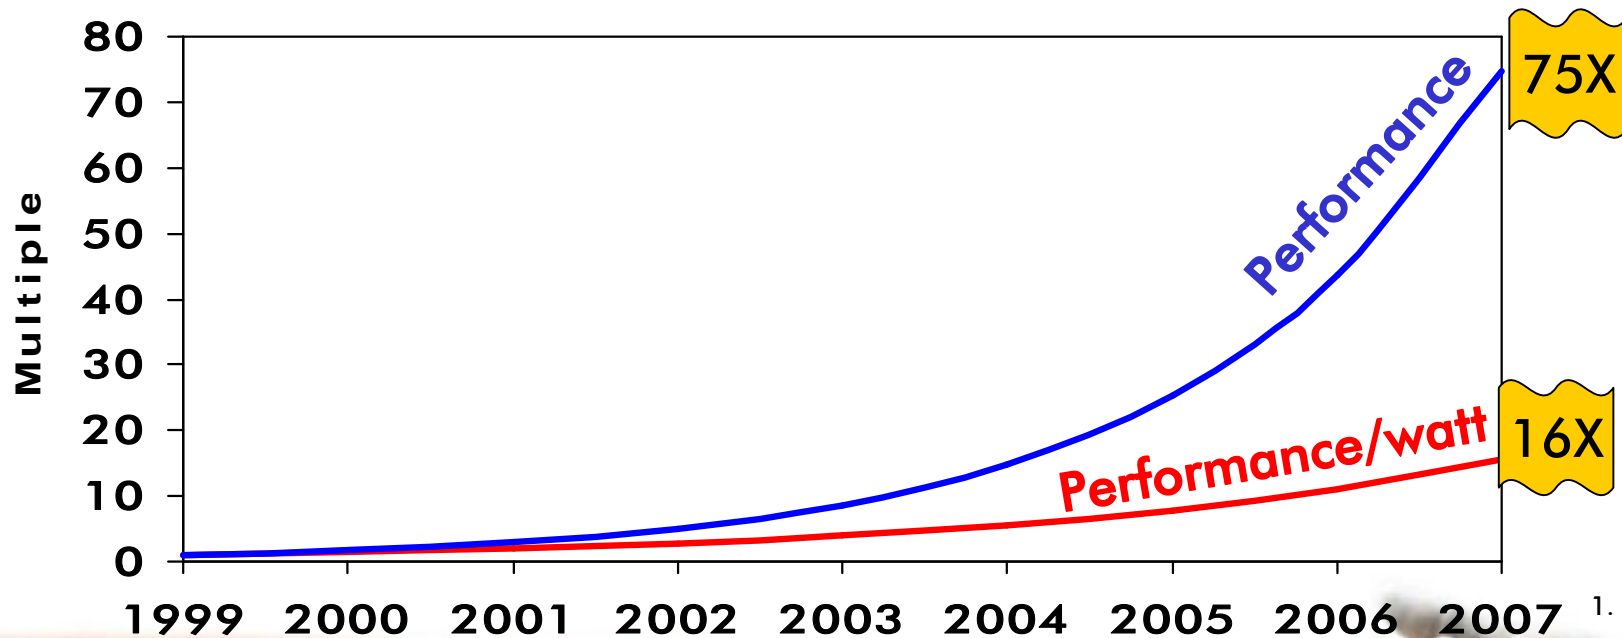
- † Approx 30% ~ 15MWhr/yr
- † 8,000 tons of CO₂
- † Equivalent to ~540 Cars



Pressure to Reduce Carbon Footprint



- t Server performance increased 75x
- t Performance/watt has not kept pace
- t Heat generation has increased 4x
- t Smaller more concentrated locations



1) Belady C., " In the data center, power and cooling costs more than the IT equipment it supports," Electronics Cooling , volume 13 no1 February 2007

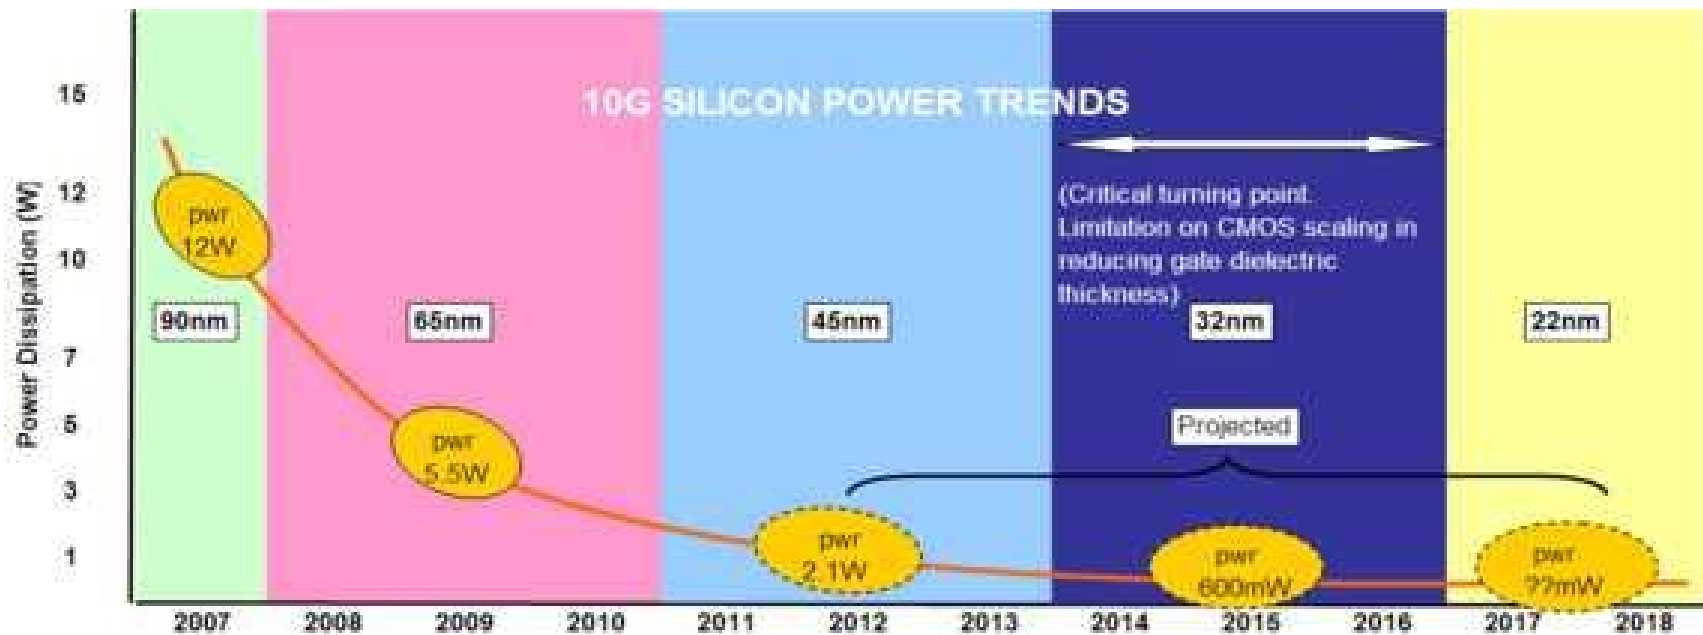
1. **Power down** servers and switches when idle
2. **Reduce** power when transceivers are in use
3. **Monitor and manage** where energy is used



Energy Efficient Ethernet (EEE)



- t Power dissipation remains a concern
- t Initial copper 10G chips consumed 12 watts
- t Move towards lower power dissipation



*Moving to next-generation process technology reduces 10GBASE-T silicon Power
(Combined Sources, ITRS, IEEE802.3)*



t Full reach at 3.5W/port




Quad 10GBASE-T Product Brief

10GBASE-T Ethernet PHY Transceiver

AQ1401 / AQ1202 / AQ1103

Aquantia's quad port 10GBASE-T PHYs provide the low power and high density to enable 48 port switches and line cards. Dual port 10GBASE-T PHYs provide the low power demanded by NICs and LOM. The single port member of Aquantia's high density family provides an upgrade path to lower power for platforms based on the 1st generation AQ1002.

These products offer low 2W/port power for up to 10m (Data Center mode) with 3.5W/port at 100m. Flexible system interfaces support XFI / RXAUI / XAUI & SGMII. A pass-through option for XFI enables dual media applications. Multi-rate capability are supported; 10G/1G/100M.

Quad – 48 port switch platforms will exploit the high port density provided with the AQ1401. Its 29x29 mm BGA package provides an excellent electrical and thermal environment. Its flexible XFI / RXAUI interface supports a wide range of high density switches

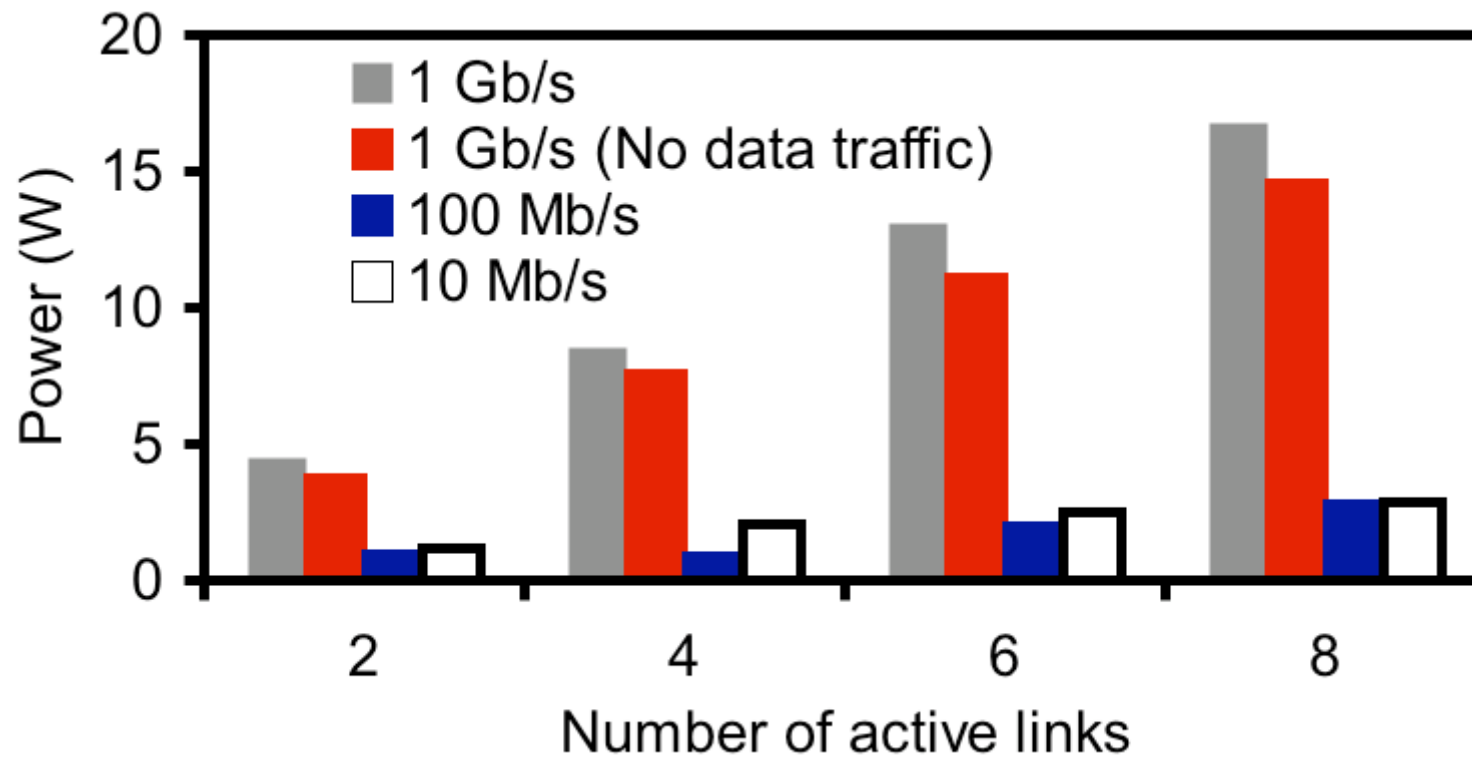
Dual – NIC and LOM applications will leverage the low power of the AQ1202 to create passively cooled solutions for the volume server market. Its 25x25 mm flip-chip BGA package brings out the full XAUI interface in addition to XFI and RXAUI.

Single – The single member of Aquantia's 2nd generation family provides backward compatibility to Aquantia's first generation AQ1002, while providing the options of XFI and RXAUI. The low power of the AQ1103 provides an excellent upgrade path for platforms based on the AQ1002.

FEATURES	SUMMARY OF BENEFITS
<ul style="list-style-type: none"> • Low Power <ul style="list-style-type: none"> - Data Center mode at 2.0W/port - Full-reach at 3.5W/port • Robust Performance <ul style="list-style-type: none"> - Exceeds BER of 10⁻¹³ over 100m of CAT5A (4-conductor mode) • Multi-rate Connectivity <ul style="list-style-type: none"> - Supports 10G/1G/100M • Advanced Cable Diagnostics <ul style="list-style-type: none"> - Enables easy connectivity debug • Built-In Thermal Management <ul style="list-style-type: none"> - On-chip thermal sensor with alarm and warning thresholds • Wake-On-LAN Support <ul style="list-style-type: none"> - Integrated circuitry to monitor for "magic-packet" • High-Performance Interfaces <ul style="list-style-type: none"> - XFI / RXAUI / XAUI / SGMII • Pass-Through Option <ul style="list-style-type: none"> - XFI to XFI pass-through option for dual media applications 	<ul style="list-style-type: none"> • Full-speed 10G Performance <ul style="list-style-type: none"> - Supports up to 100m on CAT5A - Supports up to 55m on CAT5 - Supports up to 45m on CAT5e • Familiar RJ46 and Copper Cabling <ul style="list-style-type: none"> - Delivers benefits of 10GBASE-T without the costs/complexities of SFP+ • Advanced Diagnostics <ul style="list-style-type: none"> - Enables intelligent channel monitoring, thus facilitating advanced power management • Simplified Board Design <ul style="list-style-type: none"> - Compact package reduces PCB real-estate • Dual Media Applications <ul style="list-style-type: none"> - Flexibility to support copper or fiber • Enables High-Density Switches <ul style="list-style-type: none"> - Low-power Quad-port enables high-density switches • Low EMI emissions • Flexibility in existing installations <ul style="list-style-type: none"> - 10G/1G/100M support ensures backward compatibility



Typical switch with 24 ports 10/100/1000 Mb/s

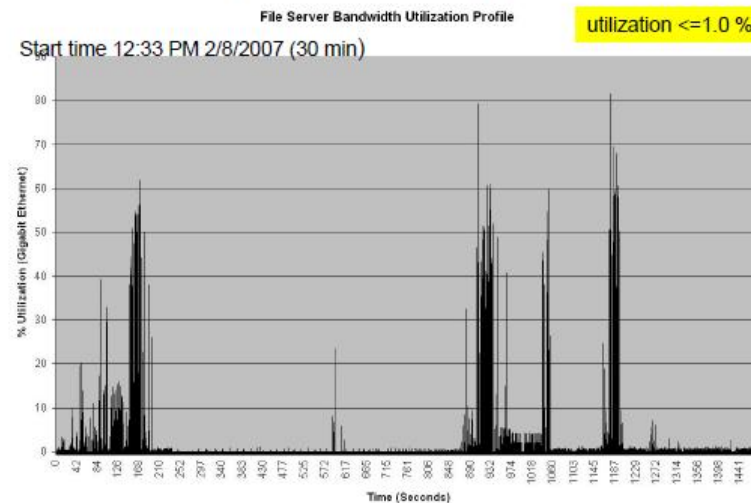


Source: Broadcom



Ethernet Traffic Profiles

- Snapshot of a File Server with 1 Gb Ethernet link
 - Shows time versus utilization (trace from LBNL)



t IT equipment consumes a lot of power in idle state

n E.g. Server ~ 80% of full load power

n Network ports in switches and NICs ~ 90% of full load power



t New set of standards (IEEE P802.3az) aimed at reducing power consumption of active Ethernet devices

n Servers, PCs, NICS and switches

n Copper only

EEE Objectives

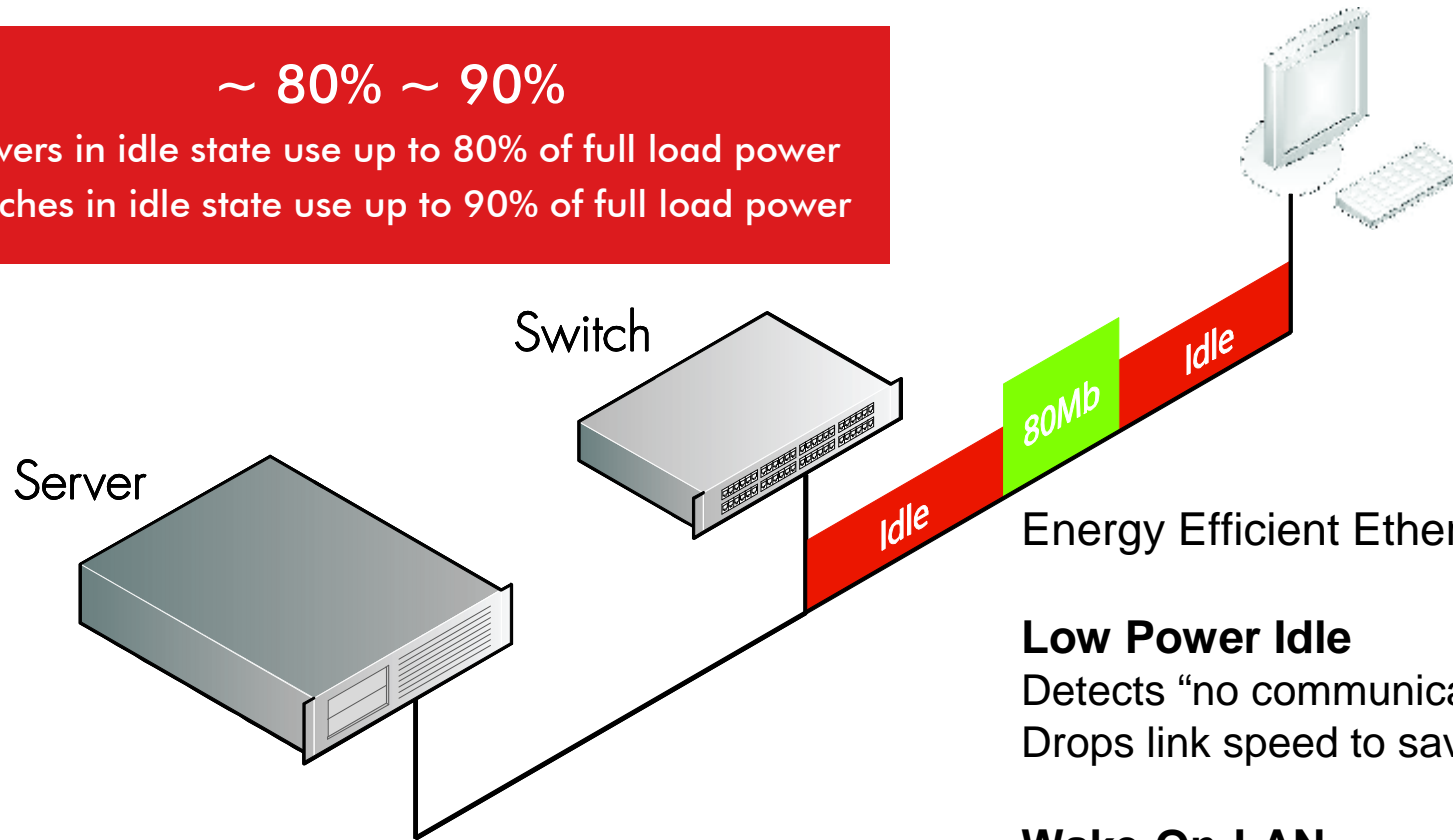
- Define a mechanism to reduce power consumption during periods of low link utilization for the following PHYs
 - 100BASE-TX (Full Duplex)
 - 1000BASE-T (Full Duplex)
 - 10GBASE-T
 - 10GBASE-KR
 - 10GBASE-KX4
 - 1000BASE-KX
- Define a protocol to coordinate transitions to or from a lower level of power consumption
- The link status should not change as a result of the transition
- No frames in transit shall be dropped or corrupted during the transition to and from the lower level of power consumption
- The transition time to and from the lower level of power consumption should be transparent to upper layer protocols and applications
- Define a 10 megabit PHY with a reduced transmit amplitude requirement such that it shall be fully interoperable with legacy 10BASE-T PHYs over 100 m of Class D (Category 5) or better cabling to enable reduced power implementations
- Any new twisted-pair and/or backplane PHY for EEE shall include legacy compatible auto negotiation



- t IEEE 802.3az Task Force initiatives
 - n Low Power Idle (LPI)
 - n Wake-on-LAN (WOL) feature



~ 80% ~ 90%
Servers in idle state use up to 80% of full load power
Switches in idle state use up to 90% of full load power



Energy Efficient Ethernet (EEE):

Low Power Idle

Detects “no communication”
Drops link speed to save power


Wake-On-LAN

Brings devices back out of sleep



t Cisco / Intel LAB Tests demonstrated 15% direct Energy Savings by enabling EEE on Cisco 4500 Switch (fully loaded with 384 ports)

- | 892 Watts without EEE
- | 751Watts with EEE



```
4510_Sup7E#show power mod
```

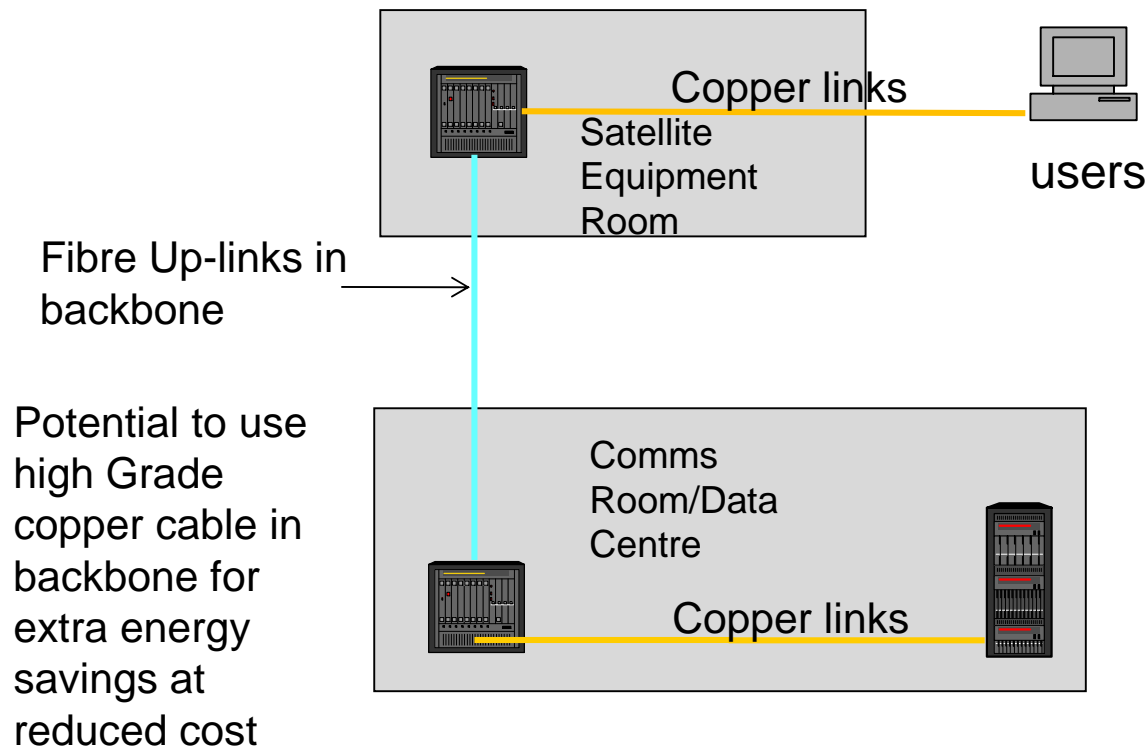
Mod	Model	Watts Used of System Power(12V)				
		budgeted	instantaneous	peak	out of reset	in reset
1	WS-X4748-UPOE+E	75	59	59	75	35
2	WS-X4748-UPOE+E	75	62	44	75	35
3	WS-X4748-MSPOE+E	75	61	62	75	35
4	WS-X4748-MSPOE+E	75	61	62	75	35
5	WS-X45-SUP7-E	260	202	205	260	100
6	WS-X45-SUP7-E	225	199	203	225	225
7	WS-X4748-MSPOE+E	75	64	64	75	35
8	WS-X4748-MSPOE+E	75	62	62	75	35
9	WS-X4748-UPOE+E	75	59	59	75	35
10	WS-X4748-MSPOE+E	75	63	63	75	35
--	Fan Tray	255	--	--	--	--
Total		1340	892	883	1085	603

```
4510_Sup7E#show power mod
```

Mod	Model	Watts Used of System Power(12V)				
		budgeted	instantaneous	peak	out of reset	in reset
1	WS-X4748-UPOE+E	75	42	59	75	35
2	WS-X4748-UPOE+E	75	44	44	75	35
3	WS-X4748-MSPOE+E	75	44	62	75	35
4	WS-X4748-MSPOE+E	75	43	62	75	35
5	WS-X45-SUP7-E	260	202	205	260	100
6	WS-X45-SUP7-E	225	199	203	225	225
7	WS-X4748-MSPOE+E	75	45	64	75	35
8	WS-X4748-MSPOE+E	75	45	62	75	35
9	WS-X4748-UPOE+E	75	42	59	75	35
10	WS-X4748-MSPOE+E	75	45	62	75	35
--	Fan Tray	255	--	--	--	--
Total		1340	751	883	1085	603

Source: Cisco / Intel White Paper : IEEE 802.3az Energy Efficient Ethernet: Build Greener Networks





Low-Power Idle (LPI):

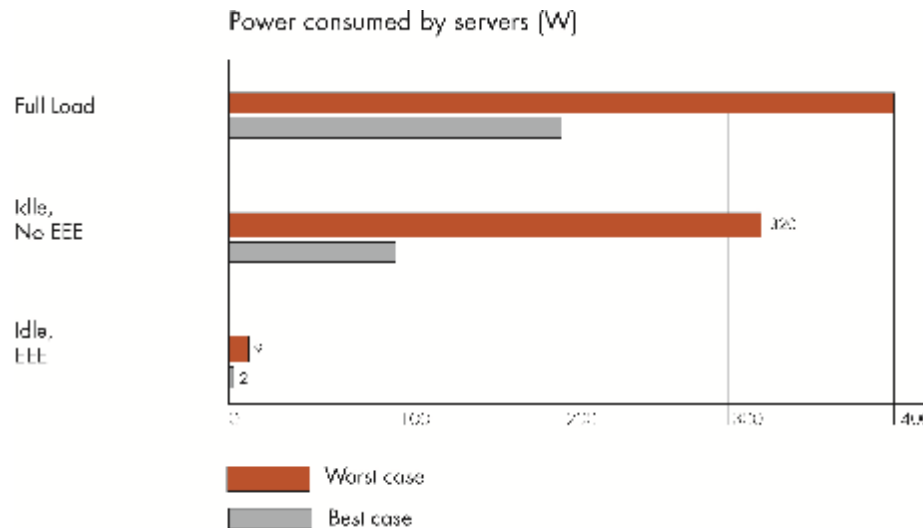
- Satellite Equipment Room switches and user PCs
- Satellite Equipment Room and Comms Room switches
- Comms Room switches and servers



t Wake-on-LAN (WOL)

- n** Servers can consume ~ 80% of full load power when idle
- n** WOL feature turns server power down when idle
- n** Assigns small percentage of standby power to network card
- n** Restarts server when 'magic packet' is received
- n** >>> Enormous power savings





Large energy savings possible

(e.g. up to 58% saving at 40% utilisation)

- Works on Ethernet auto-negotiation
- **Requires copper cabling** – auto-negotiation not available on fibre
- External noise could affect recovery time
- Shielded cabling better choice



Additional Opportunities for Further Energy Savings



t Power Backoff

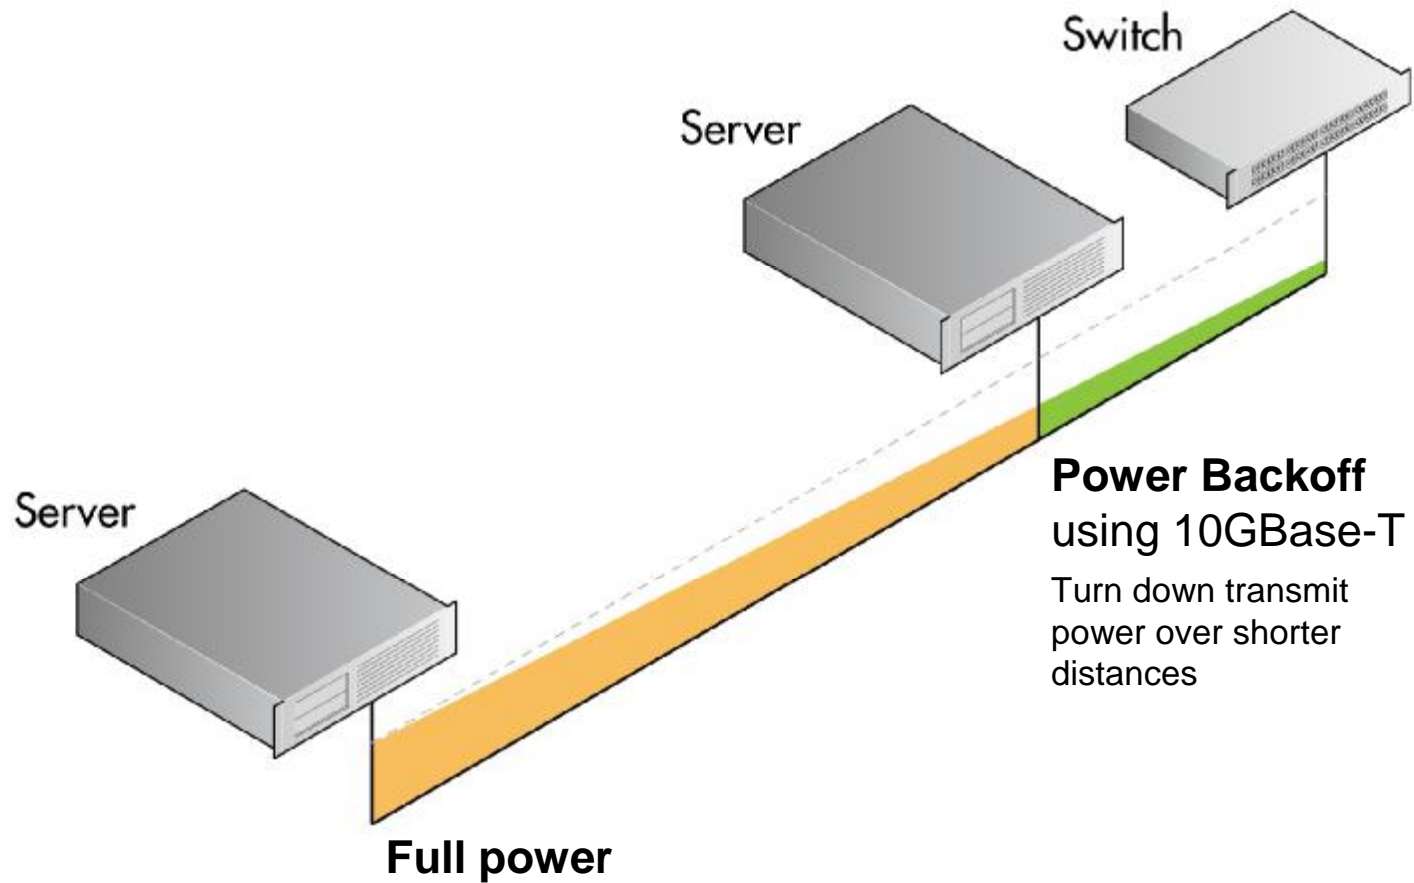
- n Mixed long and short link Alien Next problems inherent in UTP 10G cabling created requirement for Power Backoff in 802.3an standard
- n Power Backoff feature already implemented in 10GBaseT transceiver chipsets

t "Short Reach" or "Low Power" Mode of 10GBase-T

- n still under evaluation to determine power savings vs. cost of implementation



Reduce power when transceivers are in use



- t** IEEE introduced power backoff scheme to compensate for short link/long link ANEXT
- t** Transceiver establishes length of channel
- t** Transceiver turns down power for distances below 95m
- t** Scale based upon distance and measured by insertion loss

Power Backoff Scale		
Length (m)	IL 250 MHz (dB)	Backoff (dB)
0-25	<9.0	14
25-45	9.0-16.2	12
45-55	16.2-19.8	10
55-65	19.8-23.4	8
65-75	23.4-26.9	6
75-85	26.9-30.5	4
85-95	30.5-34.1	2
>95	.34.1	0



t Power Reduction of ~40% can be achieved by Power Scaling due to length

n	100m	3,5W	100%
n	10m	2 W	57%

AQUANTIA
Quad 10GBASE-T Product Brief
10GBASE-T Ethernet PHY Transceiver

AQ1401 / AQ1202 / AQ1103

Aquantia's quad port 10GBASE-T PHYs provide the low power and high density to enable 48 port switches and line cards. Dual port 10GBASE-T PHYs provide the low power demanded by NICs and LOM. The single port member of Aquantia's high density family provides an upgrade path to lower power for platforms based on the 1st generation AQ1002.

These products offer low 2W/port power for up to 10m (Data Center mode) with 3.5W/port at 100m. Flexible system interfaces support XFI / RXAUI / XAUI & SGMII. A pass-through option for XFI enables dual media applications. Multi-rate capability are supported; 10G/1G/100M.

Quad – 48 port switch platforms will exploit the high port density provided with the AQ1401. Its 29x29 mm BGA package provides an excellent electrical and thermal environment. Its flexible XFI / RXAUI interface supports a wide range of high density switches.

Dual – NIC and LOM applications will leverage the low power of the AQ1202 to create passively cooled solutions for the volume server market. Its 25x25 mm flip-chip BGA package brings out the full XAUI interface in addition to XFI and RXAUI.

Single – The single member of Aquantia's 2nd generation family provides backward compatibility to Aquantia's first generation AQ1002, while providing the options of XFI and RXAUI. The low power of the AQ1103 provides an excellent upgrade path for platforms based on the AQ1002.

Technology	Benefits
48 port switches - compact size	Low power - 3.5W/port
48 port switches - compact size	Low power - 3.5W/port
48 port switches - compact size	Low power - 3.5W/port

Aquantia Corporation, 199 Terman Drive, Milpitas, CA 95035
+1 408 255 2346
www.aquantia.com

AQUANTIA
Quad 10GBASE-T Product Brief
10GBASE-T Ethernet PHY Transceiver

AQ1401 / AQ1202 / AQ1103

Aquantia's quad port 10GBASE-T PHYs provide the low power and high density to enable 48 port switches and line cards. Dual port 10GBASE-T PHYs provide the low power demanded by NICs and LOM. The single port member of Aquantia's high density family provides an upgrade path to lower power for platforms based on the 1st generation AQ1002.

These products offer low 2W/port power for up to 10m (Data Center mode) with 3.5W/port at 100m. Flexible system interfaces support XFI / RXAUI / XAUI & SGMII. A pass-through option for XFI enables dual media applications. Multi-rate capability are supported; 10G/1G/100M.

Quad – 48 port switch platforms will exploit the high port density provided with the AQ1401. Its 29x29 mm BGA package provides an excellent electrical and thermal environment. Its flexible XFI / RXAUI interface supports a wide range of high density switches.

Dual – NIC and LOM applications will leverage the low power of the AQ1202 to create passively cooled solutions for the volume server market. Its 25x25 mm flip-chip BGA package brings out the full XAUI interface in addition to XFI and RXAUI.

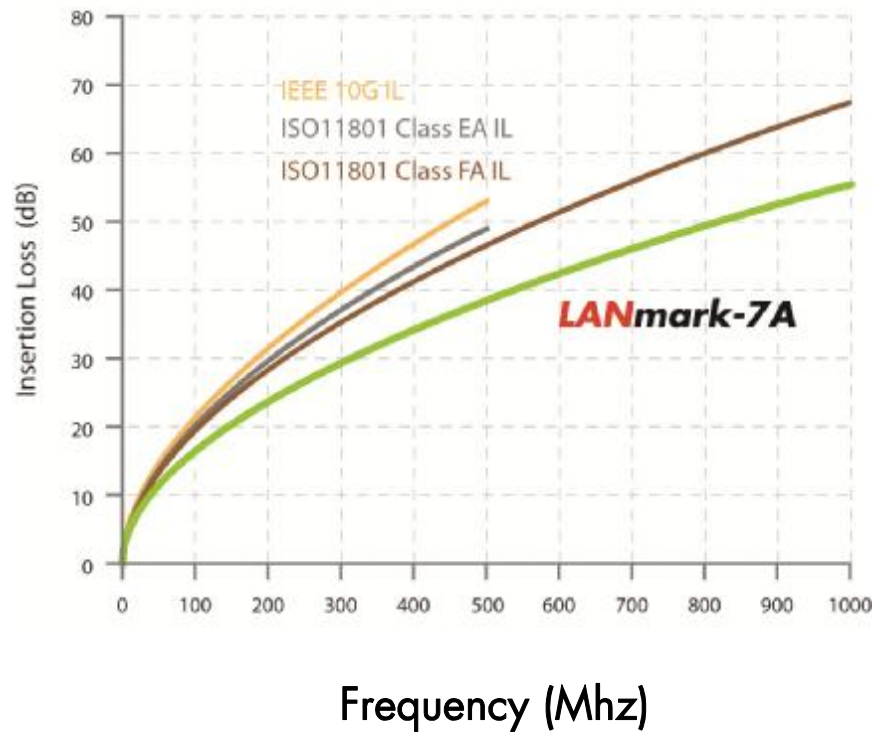
Single – The single member of Aquantia's 2nd generation family provides backward compatibility to Aquantia's first generation AQ1002, while providing the options of XFI and RXAUI. The low power of the AQ1103 provides an excellent upgrade path for platforms based on the AQ1002.

FEATURES	SUMMARY OF BENEFITS
<ul style="list-style-type: none"> Low Power <ul style="list-style-type: none"> Data Center mode at 2.0W/port Full-reach at 3.5W/port Robust Performance <ul style="list-style-type: none"> Exceeds BER of 10⁻¹² over 100m of CAT5A (4-Pair mode) Multi-rate Connectivity <ul style="list-style-type: none"> Supports 10G/1G/100M Advanced Cable Diagnostics <ul style="list-style-type: none"> Enables easy connectivity debug Built-In Thermal Management <ul style="list-style-type: none"> On-chip thermal sensor with alarm and warning thresholds Wake-On-LAN Support <ul style="list-style-type: none"> Integrated circuitry to monitor for "magic packet" High-Performance Interfaces <ul style="list-style-type: none"> XFI / RXAUI / XAUI / SGMII Pass-Through Option <ul style="list-style-type: none"> XFI to XFI pass-through option for dual media applications 	<ul style="list-style-type: none"> Full-speed 10G Performance <ul style="list-style-type: none"> Supports up to 100m on CAT5A Supports up to 55m on CAT5 Supports up to 45m on CAT5e Familiar RJ45 and Copper Cabling <ul style="list-style-type: none"> Delivers benefits of 10GBASE-T without the costs/complexities of SFP+ Advanced Diagnostics <ul style="list-style-type: none"> Enables intelligent channel monitoring, thus facilitating advanced power management Simplified Board Design <ul style="list-style-type: none"> Compact package reduces PCB real-estate Dual Media Applications <ul style="list-style-type: none"> Flexibility to support copper or fiber Enables High-Density Switches <ul style="list-style-type: none"> Low-power Quad-port enables high-density switches Low EMI emissions Flexibility in existing Installations <ul style="list-style-type: none"> 10G/1G/100M support ensures backward compatibility

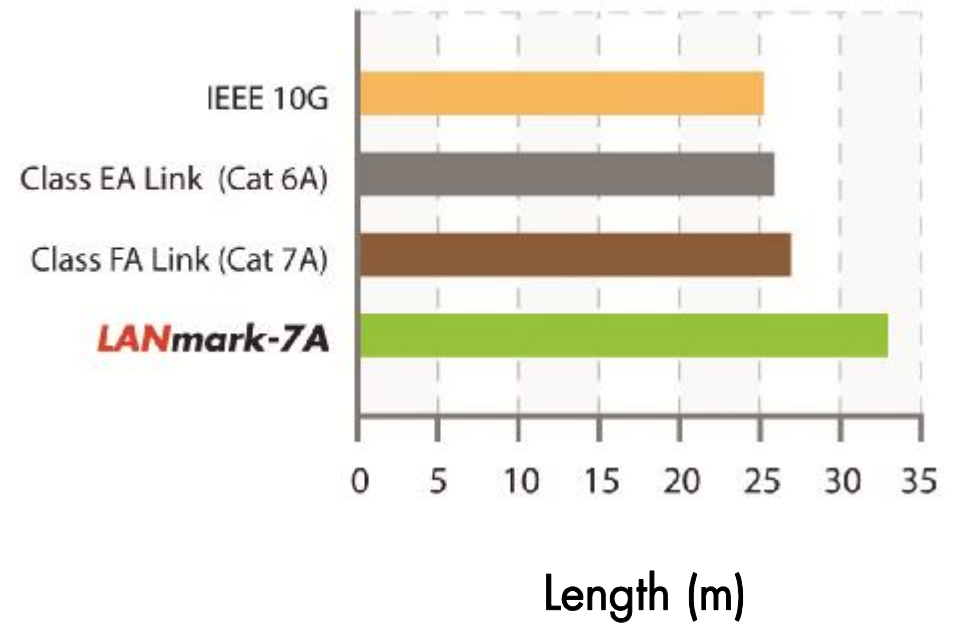
- t The Power Backoff feature makes better cabling attractive
- t Less Insertion loss can be achieved by shorter distance or higher cable grade
- t The higher the cable grade, the lower insertion loss
- t This translates into less signal strength and lower power consumption for input power = Energy Savings
- t Independantly of A-XT issues



Channel insertion loss
Cat6A / Cat7A / **LANmark-7A**



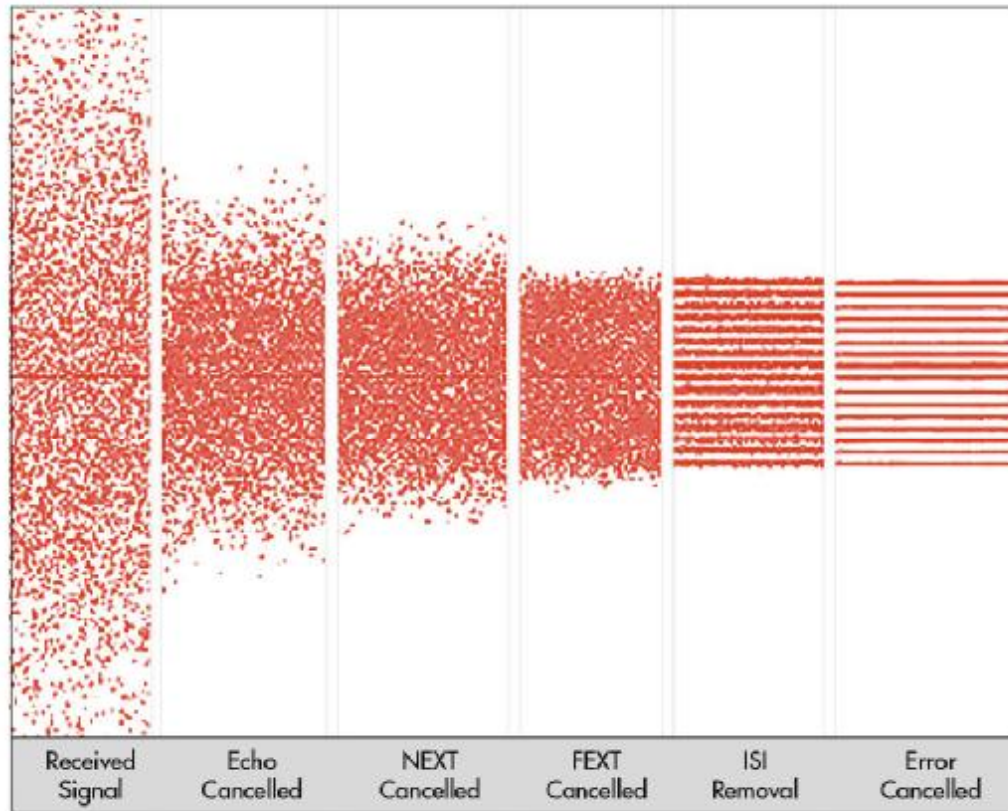
Length using Power Backoff at 9dB



- t** LANmark-7A superior Insertion Loss can extend reach of 10G for the same Power Level
- t** Long server rows more energy efficient (less CRAC required)
 - |** Server Rows can be **10 cabinets** longer using Cat.7A
 - |** Server Rows can be **140 Servers** longer in low power mode
- t** LANmark-7A can cover **90%** of all installed data centre links at full power backoff



Reduce power when transceivers are in use (2)



Source: Solarflare Communications, Inc.

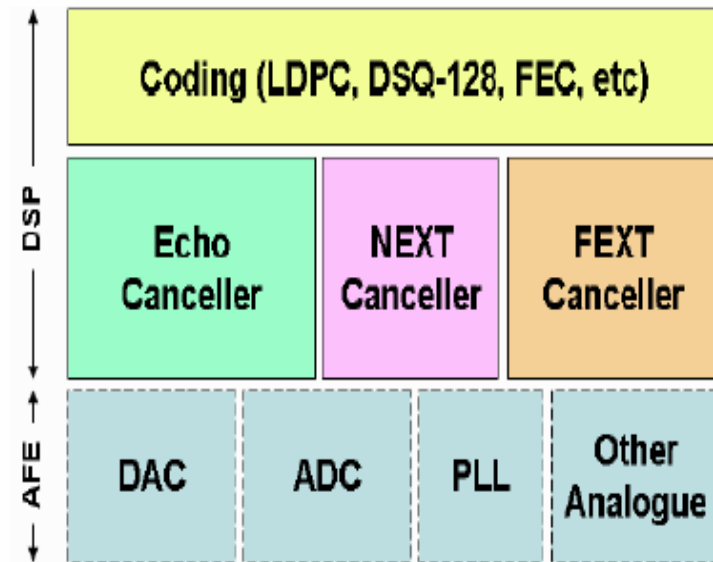
Technique	Estimated Savings
Disabling Viterbi Decode	15%
Disabling NEXT cancellers	15%
Reducing Echo canceller taps	10%
Reducing equaliser taps	5%
Reducing bit widths	5%
Potential savings	50%

Source: CommsDesign, 25th September 2009

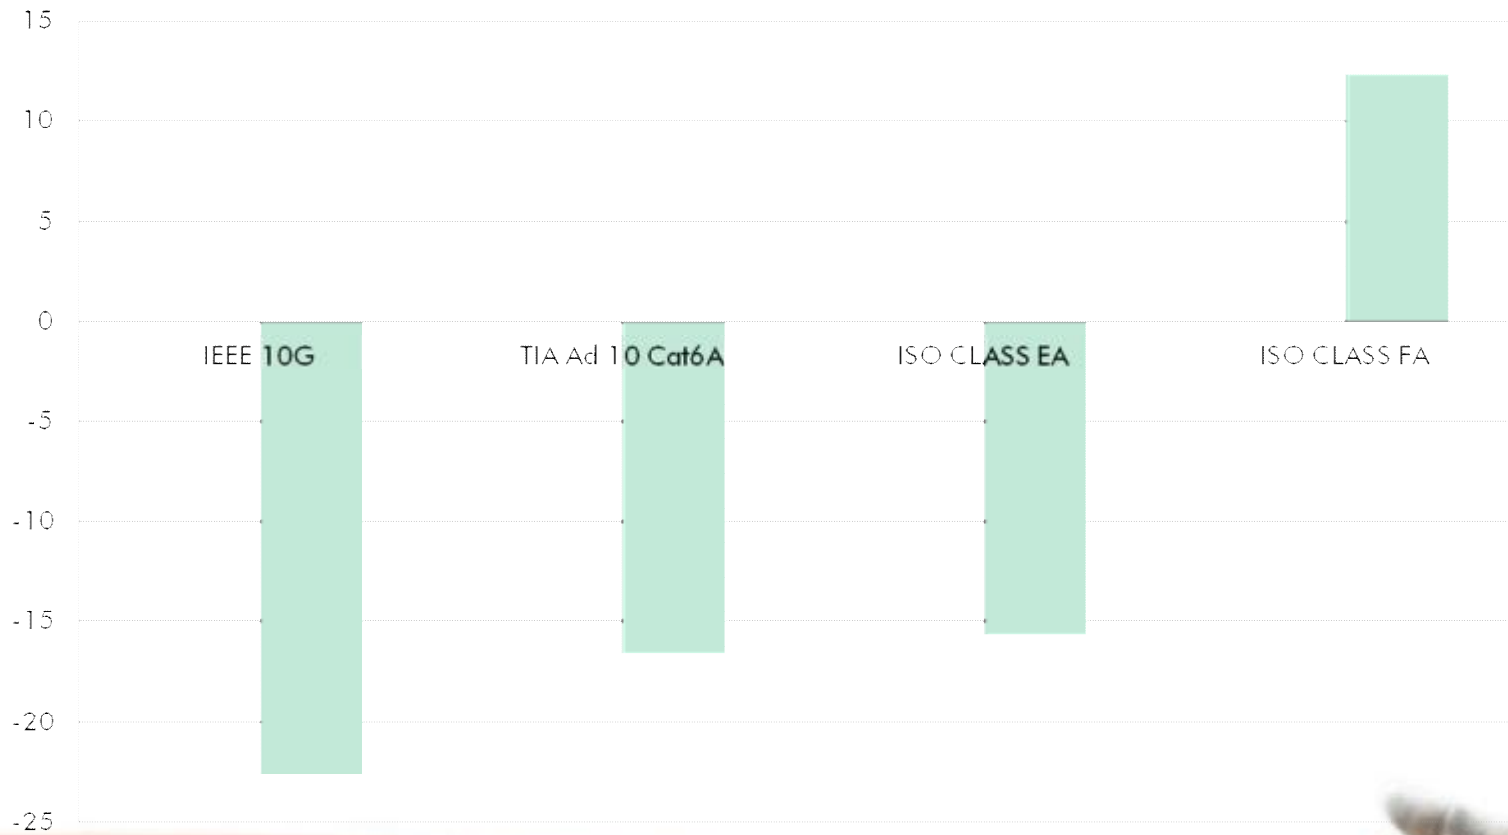


t 10GBASE-T Silicon Example

- n Echo, NEXT, and FEXT cancellers account for ~ 40% of the area and roughly 40 % of the power consumption
- n The coding elements account for ~ 25% of the area, while the remaining 35% of the chip's area is taken up for the AFEXT cancellation



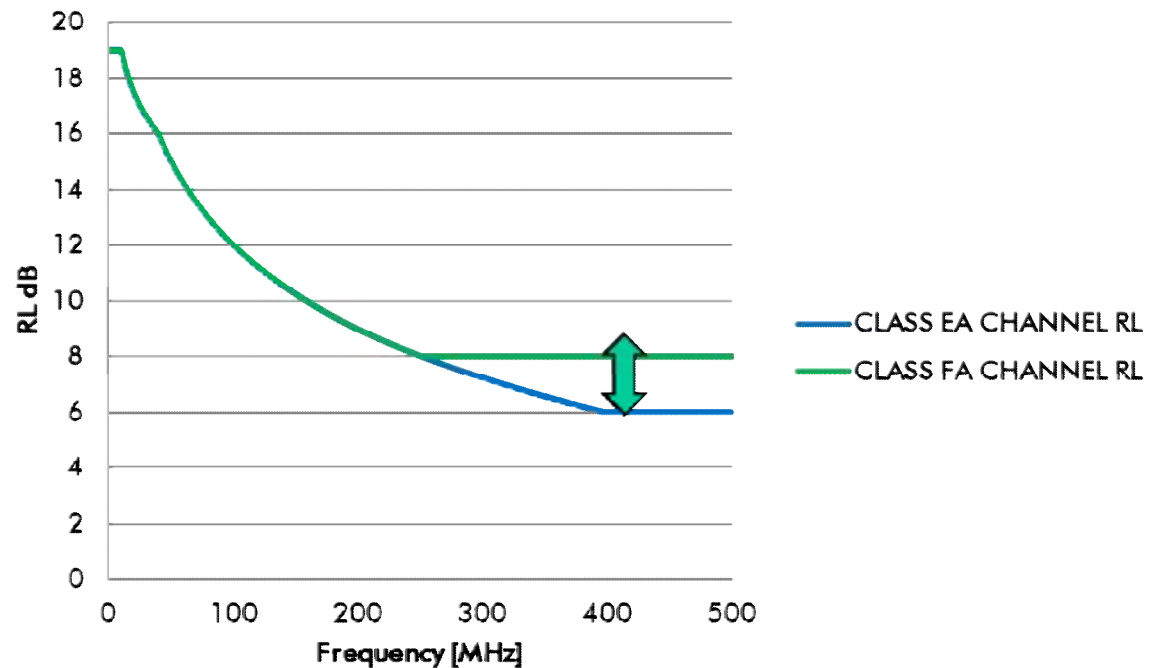
Signal to Noise Ratio of various cabling at 10G frequency (420Mhz)



Nexans Class EA versus CLASS FA RL Headroom for 10G Power Saving

- † Class FA offers 2 dB lower RL at 420 MHz
- † 25% better as limit line for Cat6A
- † Enables less ECHO Cancellation in active equipment

ISO 11801 Channel Return Loss Class EA and FA



- t Recent Publication from Aquantia announce that SNR probing is already implemented**
- t "Achieving low power in the data center with 10GBASE-T" by Kamal Dalmia and Ramin Shirani, Aquantia published 7/16/2010 9:21 AM EDT**
- t "...the 10GBASE-T PHY determines the quality of the link, cable reach, SNR, etc. The PHY, if designed for Data Center Mode, then has the **ability to turn off** any unnecessary set of cancellers, recognizing the environment to be within 10 m. As these cancellers constitute a majority of the power consumption, turning them off in this mode allows the part to shave off close to half the power. In this kind of innovative design, the part can also be **"forced"** into the Data Center Mode, allowing the system to be designed with a guaranteed reduced maximum power consumption."**



- t Better Cabling costs more money and requires higher investment
- t Do the Energy Savings provide a Return of Investment ?



Can investing in a future-proof network today lower IT costs over the next decade?

	Total installed cost	10G installed in 2013 40G installed in 2017		10G installed in 2016 40G installed in 2020	
		Est. lifetime	Annualised cost	Est. lifetime	Annualised cost
Cat. 7A S/FTP	€ 13.170	15	€ 878	15	€ 878
Cat. 6A F/UTP	€ 9.960	7	€ 1.423	10	€ 996
Cat. 6 U/UTP	€ 8.000	3	€ 2.667	6	€ 1.333

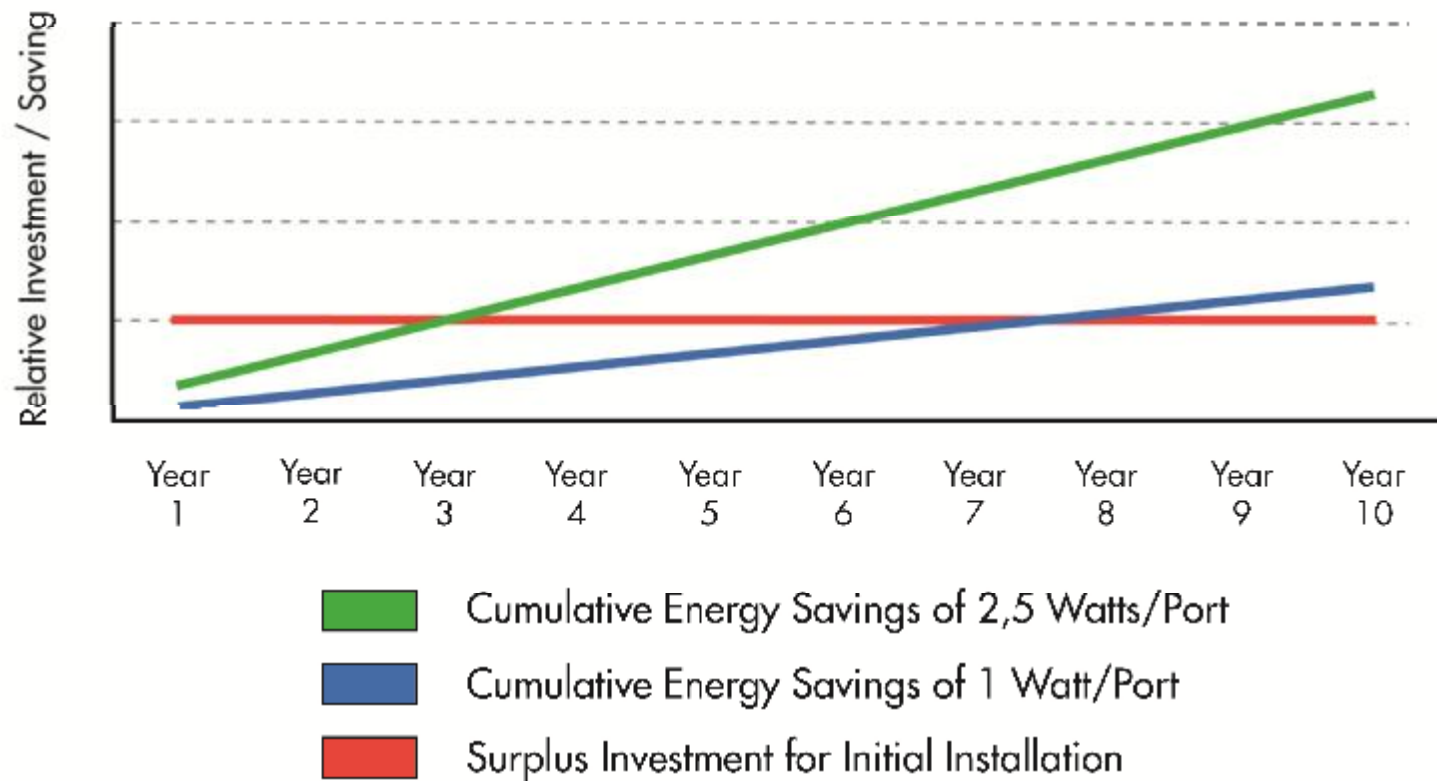
High end systems are 30% more cost efficient

Assumptions: Network of 100 channels.
Excl. costs of removal cables and downtime of network



Payback on energy saving using better cabling

Potential Power Savings of Cat.7_A compared to Cat.6_A running 10GBase-T



Additional benefits of better quality cabling



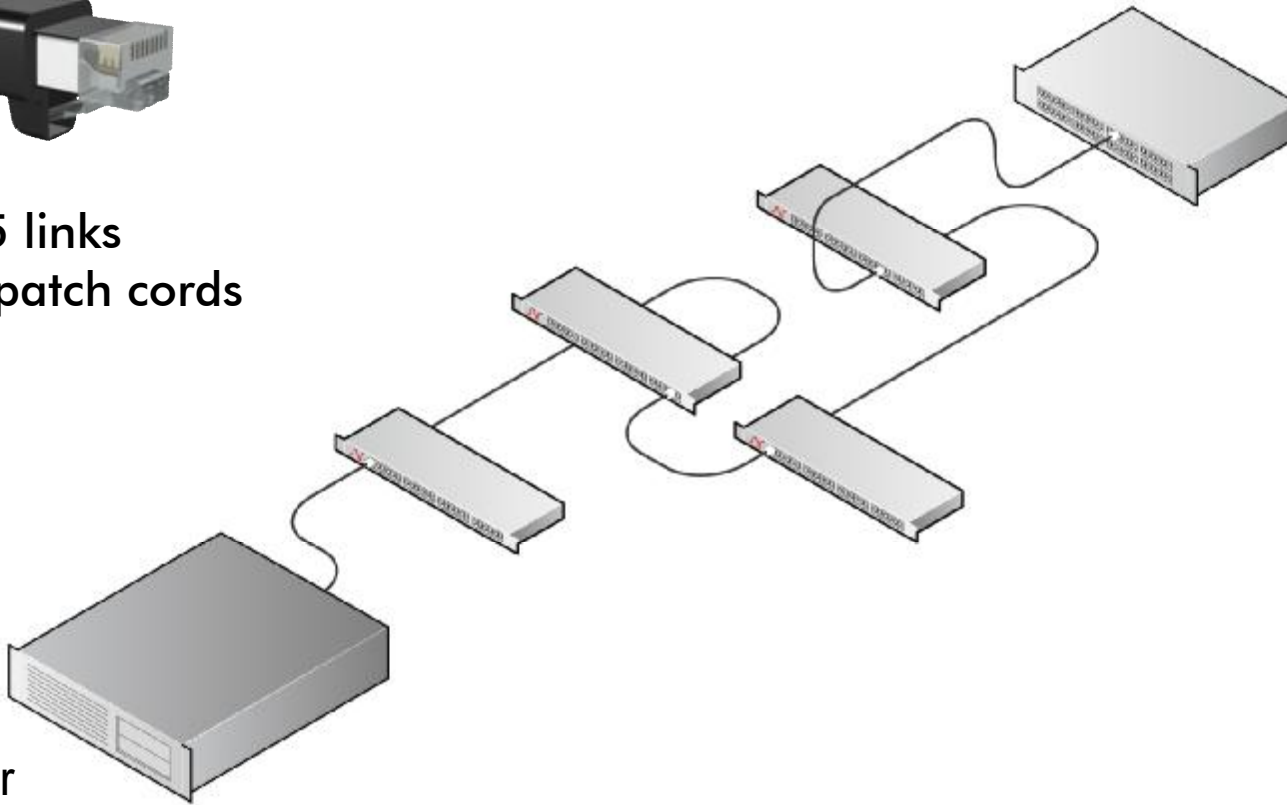


Data Centre migration path to 40G for copper





GG45 links
RJ45 patch cords




Server
10G

Switch
10G



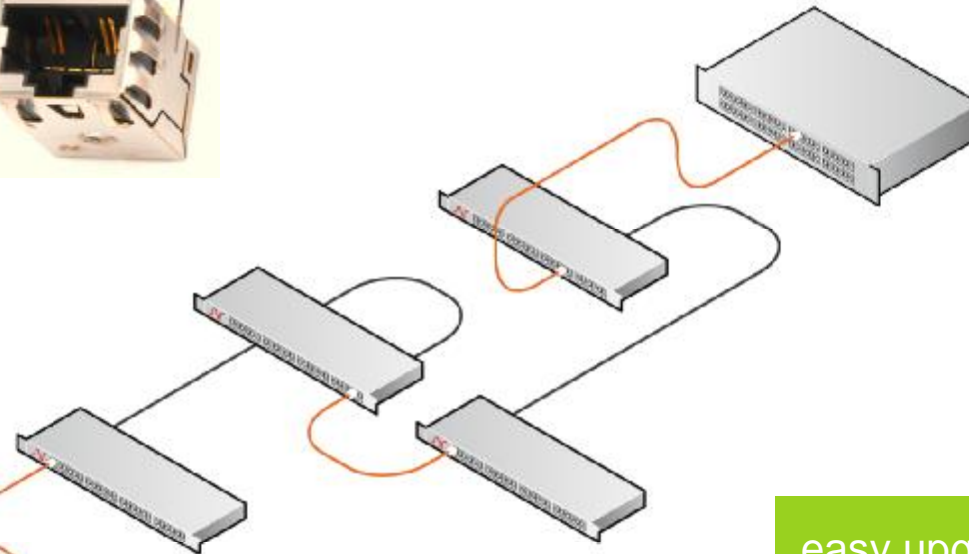
Data Centre migration path to 40G for copper



GG45 links
GG45 patch cords



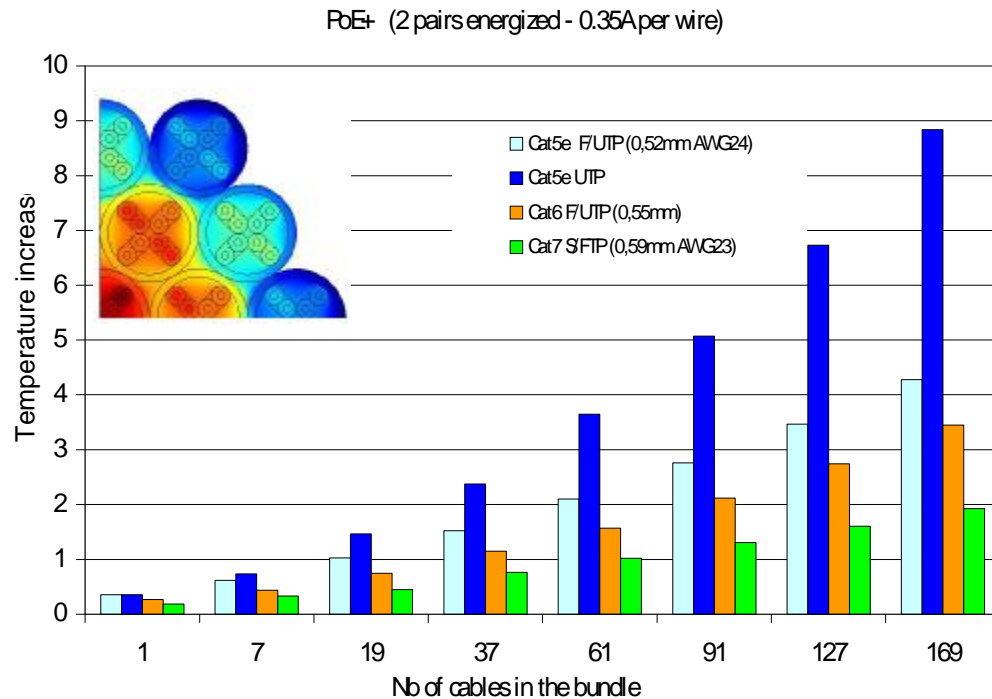
Server
40G



Switch
40G

easy upgrade
easy installation
long usable life span
low cost of ownership





- PoE boosts adoption IP convergence
- Reduces power cabling
- PoE+ supports up to 25W
- Watch effect on heat dissipation and ...
- ... on distance support and bundle size

Direct recognition from  **IEEE**
for Nexans Research Centre



- t Power Savings in IT and data centres more and more seen as mandatory not optional
- t Impact of 10Gbps NIC Cards Power Consumption not clear yet
- t IEEE802.3az encourages Copper over Fiber
- t Cat7A offers greatest potential for energy savings
 - n Enables full Power Backoff for long server rows
 - n 35%-50% less power consumption due to excess SNR / better ACR
- t 40G Migration
- t Enhance PoE support



Questions?

